

RECEIVED
CENTRAL FAX CENTERResponse dated April 6, 2007
Serial No. 09/842,604

APR 06 2007

REMARKS

Reconsideration of the rejections set forth in the Office Action is respectfully requested. Currently, claims 1-28 are pending in this application.

Rejection of claims 1-28 under 35 USC 102 over Tang

Claims 1-28 were rejected under 35 USC 102 as anticipated by Tang et al. (U.S. Patent No. 6,839,348). This rejection is respectfully traversed in view of the following arguments.

Multicast trees are commonly established to distribute information to multicast participants in an efficient manner. As noted in the background of the invention, there are several routing protocols that may be used to establish a multicast tree, two of which are Protocol Independent Multicast (PIM) and Distance Vector Multicast Routing Protocol (DVMRP). (See Specification at page 2, lines 4-6). Once the multicast tree has been established, information may be transmitted over the multicast tree. Since not every network device that is connected to the multicast tree may wish to participate in every multicast, or may not be authorized to participate in the multicast, a separate set of protocols have been developed to control membership in the multicast. One common protocol that is used to control membership in a multicast is Internet Group Management Protocol (IGMP). Thus, to summarize, multicast routing protocols such as DVMRP and PIM are used establish the multicast trees or paths through the network, and the ability to participate in multicasts taking place over the trees is controlled by IGMP or a similar membership control protocol.

Tang teaches that IGMP messages should be used to enable subscribing entities to join multicasts. (Col. 2, line 61 to Col. 3, line 6). Bridges perform IGMP snooping to determine which of their ports lead to a multicast router or to an entity subscribing to the multicast. (Col. 3, lines 7-14). Tang further teaches that MOSPF, DVMRP, or PIM may be used to install multicast routes in the network element forwarding tables for use in distributing multicast messages (Col. 3, lines 15-62).

Tang then discusses a situation where the network is partitioned into multiple VLANs. Where multiple VLANs are used on a network, traffic for the different VLANs may be identified using a VLAN ID or VLAN tag. (Tang at Col. 7, lines 16-46). One type of VLAN tag is described in the IEEE 802.1Q standard, which specifies 4095 possible VLAN designations. (Tang at Col. 7, lines 22-25). Network elements will create a separate FIB for each VLAN so

Response dated April 6, 2007
Serial No. 09/842,604

that VLAN traffic may be routed using routing information for that VLAN. In operation when a network element receives VLAN tagged traffic it will identify the FIB for that VLAN and then forward the traffic using the forwarding information stored in that FIB.

Conventionally, according to Tang, a VLAN tag was used to identify a particular VLAN. Thus, where traffic was to be multicast to more than one VLAN, it would need to be sent out more than one time – once with each VLAN tag – so that the network elements could forward the traffic onto each of the intended VLANs. (Col. 5, lines 13-22). This was the case even where the multiple VLANs were reachable over the same link. Thus, Tang proposed to use additional VLAN tags to identify groups of VLANs, which Tang refers to as MVLANs. The network elements may thus create a new FIB for the MVLAN tag, so that frames tagged using the MVLAN tag may be forwarded to the several VLANs. PIM, DVMRP, or another routing protocol may be used to install routes into the new FIB associated with the MVLAN-ID in a standard manner. In operation, messages that are to be transmitted to more than one VLAN may be tagged using a MVLAN tag for the VLAN combination, and the network elements will forward the message using the multicast routing information stored in the FIB for that MVLAN tag.

Since there are only a limited number of VLAN IDs (4095 according to the IEEE 802.1Q standard) and multiple MNDs may be assigning VLAN IDs, it is important to make sure that the several MNDs don't accidentally assign the same VLAN ID to different groups of VLANs. Tang addresses this issue by having the network administrator allocate VLAN tags to each MND which are then stored by that MND in its VLAN tag source 306. (Col. 10, lines 54-60). Alternatively, a VLAN Trunk Protocol (VTP) may be used to obtain and release VLAN designations dynamically, so that the network administrator does not need to manually configure each of the MNDs with a fixed set of available VLAN tags. (Tang at Col. 10, lines 63-67),

The Examiner has taken the position that Tang teaches a method of producing a multicast tree for an application configured to use a first multicast routing protocol, from existing protocol independent multicast routing information in a network, at least some of the protocol independent multicast routing information having been created from multicast information associated with an application configured to use a second multicast routing protocol. As support for this, the Examiner has taken the position that the protocol independent multicast routing

Response dated April 6, 2007
Serial No. 09/842,604

information is the MVLAN-ID tags, and that the "application configured to use the second multicast routing protocol" is the VLAN encapsulation or tagging protocols (VTP).

The MVLAN-ID tags are not multicast routing information, but rather are tags that are used by the network elements to identify a forwarding information base that should be used to forward a particular packet/frame. Multicast routing information is then stored in the FIBs associated with the MVLAN ID so that when a frame arrives tagged using the MVLAN ID it will be forwarded to routes that extend into all of the VLANs associated with the MVLAN ID.

The MVLAN-ID is the same type of VLAN-ID that is used to identify traffic as being associated with a single VLAN, but happens to refer to more than one VLAN. For example, Tang explains that the MVLAN-ID is allocated from the set of VLAN tags stored in the VLAN tag source 306. (Tang at Col. 13, lines 3-5 "The color-limited MVLAN-IDs encompass various subcombinations of the VLAN designations for which the respective MND is responsible.") (Tang at Col. 13, lines 11-14 "Specifically, multicast controller 302 accesses and retrieves another available VLAN designation from the VLAN tag source 306 for use as a red-limited MVLAN-ID."). Thus, the MVLAN-IDs are simply VLAN tags, i.e. IEEE 802.1Q-tags, that are used to refer to multiple VLANs, and are not protocol independent multicast routing information as asserted by the Examiner.

Additionally, Tang does not teach or suggest an application configured to use a second multicast routing protocol. VLAN encapsulation or tagging protocols are not multicast routing protocols. Tang lists several multicast routing protocols, include MOSPF, DVMRP, and PIM, at Col. 3, lines 15-18, and again at Col. 3, lines 39-46, that may be used to install routes into the network element FIBs. Notably absent from this list is any VLAN encapsulation or tagging protocol. That makes sense, since VLAN tagging protocols are used to specify how traffic is to be tagged for differential treatment on the network, not to install multicast routes into forwarding information bases of the network elements.

As support for the Examiner's position that a "VLAN encapsulation or tagging protocol" is a multicast protocol, the Examiner has cited Col. 8, lines 21-22, and col. 10, lines 63-67. For convenience, the portions of Tang that include these cited areas have been reproduced below. The surrounding text has been included to put the cited portions in context:

In Col. 8, lines 17-25, Tang states:

APR 06 2007

Response dated April 6, 2007
Serial No. 09/842,604

It should be understood that network 100 is meant for illustrative purposes only and that the present invention will operate with other, possibly far more complex, network designs. Additionally, those skilled in the art recognize that other VLAN encapsulation or tagging protocols or schemes may be utilized. Furthermore, alternative arrangements for virtually associating a set of network entities with a selected VLAN domain also exist. For example, entities may be virtually associated based on their source addresses.

This portion of Tang merely states that other ways of tagging traffic might be used instead of relying on Q-tagging. For example, source address tagging might be used. It does not teach or suggest that tagging traffic using a VLAN ID transforms the tagging protocol into a multicast routing protocol.

In Col. 10, lines 46-67, Tang states “

Once the MNDs coupled to a given VLAN region have assigned responsibility for the various VLAN domains, each MND proceeds to establish its multicast VLAN identifiers (MVLAN-IDs). First, each MND establishes a single sub-regional multicast VLAN identifier (sub-regional MVLAN-ID) that encompasses all of the VLAN domains for which the respective MND is responsible. MND 122, for example, determines that it is responsible for the red, blue and green VLAN domains of VLAN region 102. In response, multicast controller 302 accesses the VLAN tag source 306 and selects an available VLAN designation (e.g., red-blue-green) for use as its sub-regional MVLAN-ID. The VLAN tag source 306 is preferably pre-configured by the network administrator with a block of numerical identifiers that are available for selection by the MND 122 as VLAN designations. The network administrator preferably ensures that there is no overlap among the VLAN designations provided to each MND coupled to the same VLAN regions. Alternatively, the MNDs may execute an extension to one or more protocols, such as the VLAN Trunk Protocol (VTP) from Cisco Systems, Inc., in order to obtain and release VLAN designations dynamically.

This portion of Tang teaches how MVLAN IDs are assigned from the group of available VLAN IDs specified by IEEE 802.1Q. Tang further states in this section, that the VTP protocol cited by the Examiner as being a multicast protocol is actually a protocol that is configured to enable network elements to assign VLAN IDs dynamically without causing overlapping VLAN designations. VTP is therefore not a multicast routing protocol, but rather enables VLAN IDs to be allocated for VLANs on the network dynamically rather than having them statically allocated by a network administrator.

Response dated April 6, 2007
Serial No. 09/842,604

Claim 1 recites: "A method of producing a multicast tree for an application configured to use a first multicast routing protocol from existing protocol independent multicast routing information in a network, at least some of the protocol independent multicast routing information having been created from multicast information associated with an application configured to use a second multicast routing protocol, the network including a plurality of network devices including at least a plurality of routers that are members of a multicast associated with the multicast tree, a set of the routers each including a protocol independent multicast database containing the protocol independent multicast routing information." Tang does not teach or suggest anything of this nature. Rather, Tang teaches a way of using tags that span across VLAN boundaries so that a multicast routing protocol such as PIM may be used to establish a route that passes into or spans between two or more VLANs. Accordingly, applicants respectfully submit that Tang fails to anticipate claim 1 and respectfully requests that the rejection of claim 1, and those claims dependent thereon, be withdrawn. The other independent claims are likewise patentable for substantially these same reasons and, accordingly, applicants respectfully request that the rejection of these other independent claims and those claims dependent thereon also be withdrawn.

Additionally, as explained in greater detail in connection with several previous rejections, applicants are focused on enabling a network management application or other application to troubleshoot a multicast tree by reading the multicast routing information associated with the tree. After a multicast tree has been created using a particular routing protocol, it may be necessary to go back to read the routing information associated with the tree. Generally, from a routing perspective, a particular application will generally use only one of the several developed multicast routing protocols to create a multicast tree. Once a multicast tree has been established using a multicast routing protocol, the multicast tree may be read by an application that is configured to use the same routing protocol that was used to establish the multicast tree. (Specification at page 2, lines 8-9). However, applications that are configured to read multicast information established using a particular multicast routing protocol have conventionally not been able to read information about multicast trees using a routing protocol other than that particular multicast routing protocol. For example, a network management application configured to troubleshoot DVMRP trees by reading DVMRP multicast routing information

APR 06 2007

Response dated April 6, 2007
Serial No. 09/842,604

would not be able to troubleshoot a multicast tree established using PTM. (See e.g. specification at page 2, lines 9-11).

Applicant discovered that multicast information could be stored in a Management Information Base (MIB) on routers on the network in a protocol neutral format and then retrieved by applications using an available network management protocol, such as Simple Network Management Protocol (SNMP). By storing the routing information in a protocol neutral format, a network management application may read the routing information regardless of what routing protocol was used to establish the multicast tree. Tang does not teach or suggest anything of this nature. Accordingly, for this additional reason, applicants respectfully request that the rejection of claim 1 be withdrawn. The other independent claims are likewise patentable for substantially these same reasons and, accordingly, applicants respectfully request that the rejection of these other independent claims and those claims dependent thereon also be withdrawn.

Conclusion

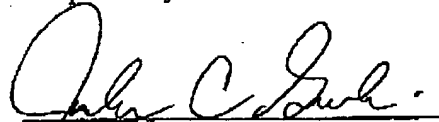
Applicants respectfully submit that the claims pending in this application are in condition for allowance and respectfully request an action to that effect. If the Examiner believes a telephone interview would further prosecution of this application, the Examiner is respectfully requested to contact the undersigned at the number indicated below.

If any fees are due in connection with this filing, the Commissioner is hereby authorized to charge payment of the fees associated with this communication or credit any overpayment to Deposit Account No. 502246 (Ref. NN-13774).

Dated: April 6, 2007

John C. Gorecki, Esq.
P.O. Box 553
Carlisle, MA 01741
Tel: (978) 371-3218
Fax: (978) 371-3219
john@gorecki.us

Respectfully Submitted


John C. Gorecki
Registration No. 38,471